Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/NA11.20260090
Published: June 9, 2026

Robust Vendor Managed Inventory Optimization with Stochastic Nonlinear Inverse Demand: A Data-Driven Swarm Intelligence Approach

Arvin Bahreini
Lundquist College of Business
University of Oregon
Eugene, Oregon, USA
arvinb@uoregon.edu

Mohammad Hossein Jalali
Graduate School of Management and Economics
Sharif University of Technology
Tehran, Tehran, Iran
mh_jalali@gsme.sharif.edu

Abstract

This paper extends the Vendor Managed Inventory (VMI) literature by integrating a robust optimization framework with stochastic nonlinear inverse demand, grounded in real-world industrial data. Unlike traditional VMI models that rely on deterministic linear demand assumptions and unconstrained financial planning, this study addresses volatile market environments where pricing power diminishes with quantity and demand is subject to random shocks. A stochastic joint replenishment problem is formulated in which the vendor maximizes a mean–variance profit objective under a strict global inventory budget. The model incorporates a quadratic inverse demand function with additive uncertainty, capturing stochastic consumer willingness to pay and diminishing marginal returns. Demand and cost parameters are estimated using the publicly available DataCo Smart Supply Chain dataset, enabling a transition from synthetic assumptions to data-driven nonlinear pricing analysis. A Simulation-Based Robust Particle Swarm Optimization (Sim-RPSO) algorithm is developed to solve the resulting non-convex and noisy optimization problem. Numerical experiments demonstrate that the proposed robust policy achieves a mean profit of $6,479 with a 95% Value-at-Risk of $6,333, outperforming the deterministic benchmark by $118.34 in mean profit and $129.37 in downside protection while reducing profit variability by 9.2% and inventory expenditure by 18.7%. These findings provide actionable managerial insights into balancing profitability and risk in budget-constrained, volatile supply chains.

Keywords
Vendor Managed Inventory (VMI), Robust Optimization, Stochastic Demand, Particle Swarm Optimization, Mean-Variance Optimization

1. Introduction
Modern supply chains operate in environments of pervasive uncertainty. Vendor Managed Inventory (VMI), a coordination mechanism in which the supplier assumes direct responsibility for replenishment at the buyer's location, has long been recognized as a powerful tool for supply chain efficiency. Pioneering work by Waller et al. (1999) and Achabal et al. (2000) demonstrated that transferring replenishment authority to the vendor yields measurable gains in service levels and inventory turnover, while Disney and Towill (2003) documented how VMI adoption substantially

dampens the bullwhip effect. However, the theoretical foundations of most VMI models remain anchored in a deterministic world — one in which demand is known, prices are fixed, and financial resources are unconstrained.

In reality, markets exhibit nonlinear price sensitivity: as a vendor supplies larger quantities to a buyer, the marginal willingness to pay declines, following a pattern more accurately captured by a quadratic inverse demand function than a linear one. Almehdawe and Mantin (2010) demonstrated that nonlinear pricing significantly alters the optimal replenishment structure, and subsequent contributions by Kim and Park (2010) and Xiao and Xu (2013) confirmed that ignoring this nonlinearity leads to systematically suboptimal decisions.

From an algorithmic perspective, the stochastic, non-convex, and budget-constrained nature of the proposed model renders traditional gradient-based methods inapplicable. Meta-heuristic approaches, particularly Particle Swarm Optimization (PSO), have demonstrated consistent success in solving such problems (Sadeghi et al., 2013; Khajehnezhad, 2018). This study develops a Simulation-Based Robust PSO (Sim-RPSO), a novel variant in which each candidate solution is evaluated through Monte Carlo simulation and personal best updates are governed by a statistically robust comparison mechanism. A further distinguishing feature is the estimation of all model parameters directly from the DataCo Smart Supply Chain dataset, encompassing over 180,000 real order records, replacing the synthetic parameter sets that dominate the VMI optimization literature. The simultaneous integration of stochastic nonlinear demand, mean-variance optimization, data-driven calibration, and Sim-RPSO constitutes the core contribution of this paper.

1.1 Objectives

The primary objective of this study is to develop a robust, data-driven optimization framework for VMI that simultaneously addresses stochastic nonlinear demand, profit risk management, and global inventory budget constraints. Specifically, this study aims to: (1) formulate a stochastic VMI model with a quadratic inverse demand function and additive demand uncertainty; (2) integrate a mean-variance profit criterion under a global inventory budget constraint; (3) develop the Sim-RPSO algorithm to solve the resulting non-convex and noisy optimization problem; (4) estimate all model parameters from the real-world DataCo Smart Supply Chain dataset; (5) conduct a rigorous comparative evaluation of the robust policy against a deterministic benchmark; and (6) extract actionable managerial insights regarding inventory allocation strategy and buyer-level profit distribution under demand uncertainty. Each objective is addressed in a dedicated section and synthesized in the conclusion.

2. Literature Review

The study of inventory control has deep roots in operations research, originating with the Economic Order Quantity framework of Harris (1913). VMI represented a significant departure from single-echelon tradition by reassigning replenishment authority to the vendor. Waller et al. (1999) and Achabal et al. (2000) provided foundational evidence of VMI efficiency gains in retail supply chains. Disney and Towill (2003) demonstrated VMI's role in dampening the bullwhip effect, and Darwish and Odah (2010) formally developed the two-echelon VMI model with a single vendor and multiple retailers. Coordination mechanisms such as consignment contracts and sales rebate structures have been shown to further enhance VMI performance (Chen et al., 2010; Wong et al., 2009). Hariga et al. (2013) added contractual storage agreements, while Kristianto et al. (2012) proposed an adaptive fuzzy VMI control system to address imprecise demand environments.

The representation of demand in VMI models has evolved substantially from fixed-demand EOQ assumptions. Almehdawe and Mantin (2010) were among the first to incorporate inverse demand into VMI, showing that the choice of demand specification significantly alters equilibrium replenishment quantities and profit distribution. Kim and Park (2010) found that profit gains from VMI coordination are amplified under nonlinear demand, and Xiao and Xu (2013) demonstrated superior outcomes from jointly coordinating price and service level decisions under price-sensitive demand. The quadratic inverse demand form $P(y) = a - by - cy^2$ adopted here is the most parsimonious specification capable of capturing both linear price sensitivity and diminishing marginal willingness to pay (Arora et al., 2010). Sajadifar and Pourghannad (2010) laid early groundwork for stochastic extensions of the VMI joint replenishment framework, and the mean-variance criterion adopted in this study draws on the portfolio optimization tradition of Markowitz (1952), which has been progressively adapted to supply chain risk management contexts.

The complexity of modern VMI models has necessitated meta-heuristic solution approaches. Nachiappan and Jawahar (2007) and Pasandideh et al. (2011) demonstrated the effectiveness of genetic algorithms for multi-product VMI with budget constraints. Nia et al. (2014) employed Ant Colony Optimization for fuzzy VMI models with shortage allowances. PSO has emerged as one of the most consistently effective approaches: Sadeghi et al. (2013) reported that tuned PSO variants outperform competing meta-heuristics on complex VMI instances, and Sadeghi et al. (2014) developed hybrid multi-objective variants for joint VMI and redundancy allocation problems. Khajehnezhad (2018) confirmed that swarm intelligence methods produce superior solutions when nonlinear cost structures are present. Despite these advances, three interconnected gaps persist: the simultaneous treatment of nonlinear pricing and stochastic demand shocks within a single VMI framework has not been attempted; the mean-variance objective has seen very limited adoption in VMI contexts; and parameter estimation in the VMI optimization literature relies almost exclusively on synthetic data. The present study fills all three gaps simultaneously. Bahreini and Jalali (2025) developed the deterministic nonlinear VMI model with PSO that directly precedes and motivates this work.

3. Methods

This section presents the mathematical formulation and solution methodology for the proposed stochastic VMI model with nonlinear inverse demand and a global inventory budget constraint. The formulation extends the deterministic nonlinear VMI framework of Bahreini and Jalali (2025) by introducing additive demand uncertainty and a mean-variance profit objective.

3.1 Model Structure and Assumptions

We consider a two-echelon supply chain consisting of a single vendor and N buyers, indexed by j = 1, 2, ..., N, each demanding M distinct products, indexed by i = 1, 2, ..., M. This structure is consistent with the multi-product, multi-buyer VMI architecture widely adopted in the operations research literature (Sadeghi et al., 2013; Cárdenas-Barrón et al., 2012). The vendor is responsible for determining the order quantity y_{ij} for product i supplied to buyer j, as well as the replenishment cycle T_j for each buyer. These constitute the decision variables of the model.

The following assumptions govern the model formulation. First, the vendor operates under a VMI arrangement in which replenishment decisions are made unilaterally by the vendor, consistent with the coordination structure described by Darwish and Odah (2010) and Waller et al. (1999). Second, demand is continuous and realized at the time of replenishment. Third, shortages are not permitted, and all demand must be fulfilled within each replenishment cycle. Fourth, the holding cost is incurred on the average inventory level within each cycle, following the standard EOQ-based inventory cost structure (Harris, 1913). Fifth, the vendor faces a global budget constraint on total inventory-related expenditure, as introduced into the VMI context by Cárdenas-Barrón et al. (2012). A summary of all notation used in the model is provided in Appendix A.

3.2 Stochastic Nonlinear Inverse Demand Function

The inverse demand function faced by the vendor at buyer j for product i is specified as a quadratic function of the order quantity with an additive stochastic shock:

$$P_{ij}\left(y_{ij}, \varepsilon_{ij}\right) = \alpha_{ij} - \beta_{ij} y_{ij} - \gamma_{ij} y_{ij}^2 + \varepsilon_{ij}$$

where $\alpha_{ij} > 0$ is the market intercept representing the maximum willingness to pay when quantity approaches zero, $\beta_{ij} > 0$ is the linear price sensitivity coefficient, $\gamma_{ij} > 0$ is the quadratic curvature coefficient capturing diminishing marginal willingness to pay, and ε_{ij} is an additive demand shock distributed as $\varepsilon_{ij} \sim \mathcal{N}\left(0, \sigma_{ij}^2\right)$. The realized price is bounded below by zero, so that $P_{ij} = \max\left(P_{ij}\left(y_{ij}, \varepsilon_{ij}\right), 0\right)$.

The quadratic inverse demand specification has been widely employed in the VMI and supply chain coordination literature to capture the nonlinear relationship between supply quantity and market price (Almehdawe and Mantin, 2010; Xiao and Xu, 2013). Bahreini and Jalali (2025) demonstrated in a deterministic setting that this specification produces significantly higher optimal profits than a linear approximation, providing empirical justification for the additional modeling complexity. The present study extends that formulation by introducing the stochastic term ε_{ij}, which captures the random component of consumer willingness to pay arising from unobserved market conditions, seasonal fluctuations, and competitive dynamics (Sajadifar and Pourghannad, 2010).

3.3 Mean-Variance Objective Function

The vendor's optimization problem is formulated as the maximization of a mean-variance utility function over the stochastic profit distribution. The total profit realized under a given replenishment policy (y_{ij}, T_j) and demand shock realization ε is:

$$Z(y,T,\varepsilon) = \sum_{j=1}^{N}\sum_{i=1}^{M}\left[P_{ij}(y_{ij},\varepsilon_{ij})\cdot y_{ij} - \delta_{ij}y_{ij} - h_i\left(\frac{y_{ij}T_j}{2}\right)\right] - \sum_{j=1}^{N}\frac{K_j}{T_j}$$

where δ_{ij} is the unit production cost for product i at buyer j, h_i is the unit holding cost for product i per unit time, and K_j is the fixed ordering cost for buyer j per replenishment cycle. The three cost terms represent production cost, inventory holding cost, and ordering cost respectively, consistent with the joint replenishment cost structure used in the VMI literature (Khajehnezhad, 2018; Darwish and Odah, 2010).

The mean-variance objective is defined as:

$$\max_{y,T} \mathbb{E}[Z] - \lambda \cdot \mathrm{Var}[Z]$$

where $\lambda \geq 0$ is a risk aversion coefficient that governs the trade-off between expected profit and profit volatility. When λ = 0, the objective reduces to pure expected profit maximization under simulated demand. The deterministic benchmark goes one step further, replacing the expectation with the profit evaluated at ε = 0 and solving the resulting noiseless problem with standard PSO. When $\lambda > 0$, the vendor explicitly penalizes profit variance, producing a robust policy that performs more consistently across demand realizations. The mean-variance criterion is well established in portfolio optimization (Markowitz, 1952) and has been applied to supply chain risk management contexts where profit stability is a managerial priority (Kim and Park, 2010). Since $\mathbb{E}[Z]$ and $\mathrm{Var}[Z]$ are not available in closed form under the stochastic nonlinear demand specification, both quantities are estimated via Monte Carlo simulation during optimization, as described in Section 3.4.

3.4 Constraint Set

The model is subject to the following constraints. Order quantities are bounded below by zero and above by a product-specific maximum:

$$0 < y_{ij} \leq y_{ij}^{\max} \ \forall i, j$$

Replenishment cycles are bounded within a practically meaningful range:

$$T^{\min} \leq T_j \leq T^{\max} \ \forall j$$

The global inventory budget constraint limits total inventory-related expenditure across all buyers:

$$\sum_{j=1}^{N}\left(\frac{K_j}{T_j} + \sum_{i=1}^{M} h_i \frac{y_{ij}T_j}{2}\right) \leq B$$

where B is the total available budget for inventory costs. This constraint, introduced in the VMI context by Cárdenas-Barrón et al. (2012), forces the vendor to make value-based allocation decisions, directing inventory resources toward the buyer-product combinations that generate the highest risk-adjusted profit contribution. The budget constraint is enforced through a penalty function in the Sim-RPSO algorithm, with a sufficiently large penalty coefficient Λto ensure that infeasible solutions are systematically excluded from the optimal solution set (Pasandideh et al., 2011).

3.5 Simulation-Based Robust Particle Swarm Optimization (Sim-RPSO)

To solve the resulting non-convex, noisy, and constrained optimization problem, this study develops a Simulation-Based Robust Particle Swarm Optimization algorithm. Standard PSO, originally proposed by Kennedy and Eberhart (1995), maintains a population of candidate solutions, referred to as particles, each characterized by a position vector representing a candidate replenishment policy and a velocity vector governing its movement through the solution space. The position and velocity of particle k at iteration t are updated according to:

$$v_{kd}(t+1) = w\cdot v_{kd}(t) + c_1 r_1\big(p_{kd} - x_{kd}(t)\big) + c_2 r_2\big(g_d - x_{kd}(t)\big)$$
$$x_{kd}(t+1) = x_{kd}(t) + v_{kd}(t+1)$$

where w is the inertia weight, c_1 and c_2 are the cognitive and social acceleration coefficients, r_1 and r_2 are uniform random variables on [0,1], p_{kd} is the personal best position of particle k in dimension d, and g_d is the global best position across all particles. PSO has demonstrated consistent effectiveness on nonlinear VMI problems (Sadeghi et al., 2013; Khajehnezhad, 2018).

Sim-RPSO introduces two key modifications to standard PSO. First, the inertia weight w decays linearly from w_{start} to w_{end}, promoting broad exploration in early iterations and fine-grained exploitation as convergence is approached (Sadeghi et al., 2014). Second, the robust personal best update rule requires any candidate improvement to be confirmed under a higher-fidelity re-evaluation: when the current position appears to outperform the personal best under n_{eval} simulations, both are re-evaluated under $n_{robust} > n_{eval}$ simulations, and the update is accepted only if the improvement is confirmed. Both re-evaluations are performed under common random numbers, so that the comparison reflects genuine differences between the candidate and incumbent policies rather than differences in the sampled demand scenarios. This mechanism prevents the algorithm from accepting spurious improvements caused by favorable random draws in Monte Carlo evaluation, producing a more stable convergence trajectory and a more reliable final solution.

Each particle represents a candidate replenishment policy encoded as:

$$\mathbf{x}_k = [y_{11}, y_{12}, \dots, y_{MN}, T_1, T_2, \dots, T_N]$$

The fitness of each particle is evaluated as:

$$\text{Fitness}(\mathbf{x}_k) = -\mathbb{E}[Z] + \lambda \cdot \text{Var}[Z] + \Lambda \cdot \max(0, \text{InvCost}(y, T) - B)$$

where $\mathbb{E}[Z]$ and $\text{Var}[Z]$ are estimated from n_{eval} Monte Carlo simulations, and Λ is a large penalty coefficient applied for budget violations. The algorithm parameters used in this study are described in Section 5.1.

The computational complexity of Sim-RPSO scales as $O(n_{\text{particles}} \times n_{\text{iter}} \times n_{\text{sim}})$, where n_{sim} is the number of Monte Carlo simulations per fitness evaluation. While this is substantially higher than the cost of deterministic PSO, the additional computational burden is justified by the richer information about the profit distribution that simulation-based evaluation provides, information that is unavailable to deterministic solvers by construction.

4. Data Collection

The empirical analysis uses the DataCo Smart Supply Chain dataset, a publicly available transaction-level dataset published on the Mendeley Data repository (https://data.mendeley.com/datasets/8gx2fvg2k6/5) and freely accessible for academic research purposes. The dataset contains 180,519 order records spanning multiple product categories, customer segments, geographic markets, and shipping modes, representing a realistic multi-product, multi-segment supply chain environment of the kind VMI frameworks are designed to coordinate. Records are screened for missing or non-positive quantity, price, and cost values following standard data preparation practice (Arora et al., 2010), and the data is restricted to the four highest-volume product categories — Cleats, Men's Footwear, Women's Apparel, and Indoor/Outdoor Games — serving as M = 4 products, and the three customer segments — Consumer, Corporate, and Home Office — serving as N = 3 buyers, yielding 87,130 transaction records for parameter calibration and 15 decision variables, consistent with multi-product VMI studies (Sadeghi et al., 2013; Nia et al., 2014).

All parameters are estimated directly from the filtered dataset. Quadratic inverse demand parameters are calibrated via moment-matching, with sanity checks confirming predicted prices match observed means within acceptable bounds (Almehdawe and Mantin, 2010). Unit production cost is set at 40% of median list price, consistent with the cost structure employed by Arora et al. (2010) and within the 35%–45% production-to-price ratio range reported for multi-category consumer goods supply chains (Nachiappan and Jawahar, 2007; Pasandideh et al., 2011). Holding cost is set at 20% of mean unit production cost per year, a standard assumption grounded in the EOQ tradition (Harris, 1913) and widely adopted in VMI studies operating under similar cost structures (Cárdenas-Barrón et al., 2012; Sadeghi et al., 2013). Ordering cost is set at 2% of mean segment revenue per cycle, following the proportional cost specification of Darwish and Odah (2010) and consistent with the ordering cost ratios reported by Nia et al. (2014) for multi-product VMI instances. Demand shock standard deviation is the residual standard deviation of observed prices around the fitted quadratic curve, with a minimum floor of 0.5 to prevent degenerate near-zero uncertainty (Sajadifar and Pourghannad, 2010). The global budget is fixed at B = \$300, approximately 2.4 times the unconstrained optimal inventory expenditure, providing realistic financial headroom while preserving the economic role of the constraint (Cárdenas-Barrón et al., 2012); a sensitivity analysis in Section 5.3 maps the regimes in which the budget becomes binding and quantifies its effect on profitability. Bahreini and Jalali (2025) identified real-world data calibration as a critical direction for future research, and the present study responds directly to that call. Market intercepts range from \$54.98 to \$142.99, and unit production costs range from \$19.99 to \$52.00. Cleats exhibit by far the highest price sensitivity (linear coefficients of 1.92–2.36, against 0.05–0.59 for the remaining categories) and the highest demand shock uncertainty (σ between 9.46 and 11.58, against 1.53–1.57 for Women's Apparel and the 0.5

floor for the remaining categories), making this category the natural locus of risk-driven differences between robust and deterministic policies.

5. Results and Discussion

This section presents a comprehensive analysis of the numerical and graphical results obtained from applying the proposed Sim-RPSO algorithm to the stochastic VMI model calibrated on the DataCo Smart Supply Chain dataset.

5.1 Numerical Results

The Sim-RPSO algorithm was executed with 40 particles over 150 iterations, using 200 Monte Carlo simulations per standard fitness evaluation, 400 simulations for robust personal best confirmation, inertia weight decaying from 0.9 to 0.4, cognitive and social coefficients c1 = c2 = 2.0, risk aversion coefficient λ = 0.01, and budget B = $300; the deterministic benchmark was solved using standard PSO with identical population size and iteration count, and all experiments were implemented in Python 3 using the NumPy and Pandas libraries (Bahreini and Jalali, 2025). To account for the stochasticity of the search process, the optimization was repeated across five independent random seeds, with the reported robust policy being the best solution identified under high-fidelity re-evaluation, following standard multi-start practice (full run-to-run results appear in Section 5.4); a complete run requires approximately 109 seconds. The resulting optimal replenishment cycles are T_Consumer = 0.1406, T_Corporate = 0.1777, and T_Home Office = 0.1477 years — approximately every 7.3, 9.2, and 7.7 weeks respectively — whereas the deterministic benchmark drives the Home Office cycle to its lower bound of 0.05 years (2.6 weeks), incurring substantially higher fixed ordering-cost intensity. By selecting moderate cycles across all three buyers, the robust policy reduces total ordering expenditure, a structural difference that contributes directly to its lower overall inventory cost.

Table 1. Optimal Order Quantities under Robust Policy (Sim-RPSO), in units.

Product	Consumer	Corporate	Home Office
Cleats	5.039	3.814	4.493
Men's Footwear	15.000	15.000	15.000
Women’s Apparel	9.215	9.571	9.878
Indoor/outdoor Games	15.000	15.000	15.000

Table 2. Optimal Order Quantities under Deterministic Policy (Std PSO), in units.

Product	Consumer	Corporate	Home Office
Cleats	5.358	4.812	4.858
Men's Footwear	15.000	15.000	15.000
Women’s Apparel	9.644	15.000	9.596
Indoor/outdoor Games	15.000	15.000	15.000

Tables 1 and 2 present the optimal order quantities under the robust and deterministic policies respectively. In both, Men's Footwear and Indoor/Outdoor Games receive the maximum allocation of 15 units across all segments, reflecting their high margins and low demand uncertainty. The two policies diverge precisely where uncertainty is concentrated. First, the robust policy assigns uniformly smaller Cleats quantities (3.81–5.04 units against 4.81–5.36 under the deterministic policy); Cleats carry by far the largest demand shocks (σ between 9.46 and 11.58), so each unit contributes disproportionately to profit variance, and the mean-variance objective trims exposure accordingly. Second, the deterministic policy assigns the corner value of 15.000 units of Women's Apparel to the Corporate segment, while the robust policy selects an interior allocation of 9.571 units, consistent with its allocations to the other two segments. The confirmed-improvement update rule of Sim-RPSO, which accepts changes only when they survive high-fidelity re-evaluation, acts as an implicit regularizer against such extreme allocations in the noisy fitness landscape.

Table 3. Head-to-Head Performance Comparison: Sim-RPSO vs Deterministic Benchmark.

Metric	Sim-RPSO (Robust)	Std PSO (Deterministic)
Mean Profit ($)	6,479.43	6,361.09
Std Dev of Profit ($)	88.50	97.44
VaR 95% ($)	6,333.49	6,204.12
Inventory Cost ($)	123.97	152.49
Budget Utilization (%)	41.3%	50.8%

In this study, VaR at the 95% confidence level refers to the 5th percentile of the simulated profit distribution — the minimum profit the vendor achieves in 95% of demand scenarios — where a higher value indicates better downside protection (Rockafellar and Uryasev, 2000). Table 3 presents the head-to-head performance comparison. The robust policy dominates the deterministic benchmark across all five metrics simultaneously: mean profit is $118.34 higher, profit standard deviation is 9.2% lower, and the 95% VaR improves by $129.37, meaning the vendor earns at least $129 more in the worst five percent of demand scenarios. Notably, the robust policy achieves this while spending 18.7% less on inventory ($123.97 against $152.49), indicating that its advantage stems from a structurally better allocation — moderate cycles, trimmed exposure to the most volatile product, and avoidance of corner allocations — rather than from greater resource consumption.

5.2 Graphical Results

The graphical analysis draws on four figures generated from the optimization results, each providing a distinct perspective on the behavior and performance of the proposed framework.

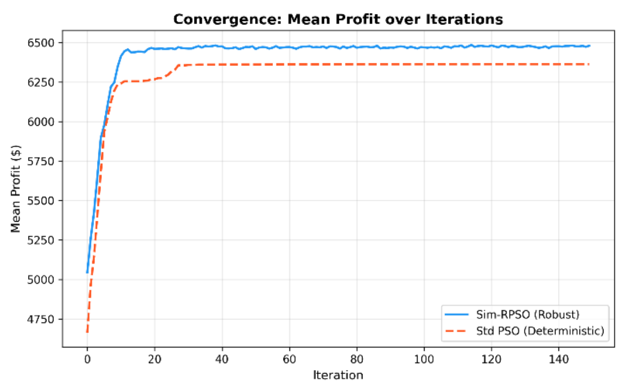

Figure 1. Convergence of Mean Profit over Iterations: Sim-RPSO vs. Standard PSO.

Figure 1 presents the convergence trajectories of both algorithms in terms of mean profit over 150 iterations. The Sim-RPSO algorithm achieves rapid initial improvement, advancing from a mean profit of approximately $5,040 at the first iteration to its plateau near $6,480 within the first 15 iterations. The standard PSO deterministic benchmark follows a similar qualitative trajectory but converges to a lower plateau of approximately $6,363 by iteration 30, reflecting the absence of the robust re-evaluation mechanism that enables Sim-RPSO to reject spurious improvements. The convergence behavior of both algorithms is consistent with the PSO convergence patterns reported for complex VMI optimization instances in the literature (Sadeghi et al., 2013; Khajehnezhad, 2018). The separation between the two convergence curves, which is fully established by iteration 30 and maintained through iteration 150, provides visual confirmation that the performance advantage of Sim-RPSO is structural rather than incidental.

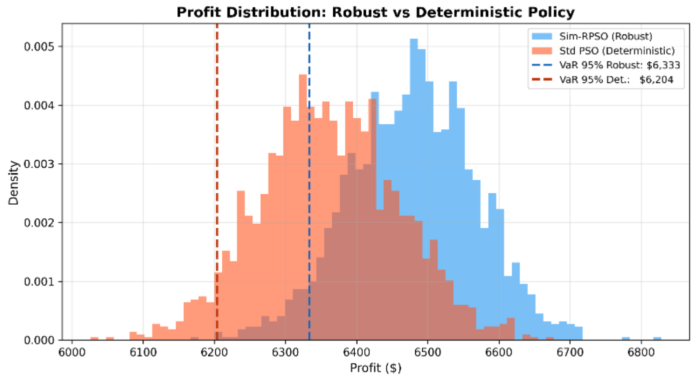

Figure 2. Profit Distribution Comparison: Robust Policy (Sim-RPSO) vs. Deterministic Policy (Std PSO).

Figure 2 presents the full profit distributions of both policies, estimated from 2,000 Monte Carlo simulations of the final optimal solutions. The robust policy distribution, shown in blue, is shifted rightward relative to the deterministic distribution, shown in orange, confirming the mean profit advantage of the robust solution. The robust distribution also exhibits a tighter concentration of probability mass around its center, consistent with the lower standard deviation reported in Table 3. The two dashed vertical lines indicate the respective 95th percentile VaR values: $6,333 for the robust policy and $6,204 for the deterministic policy. The leftward tail of the deterministic distribution extends further than that of the robust distribution, reflecting the greater downside risk exposure of the deterministic solution under adverse demand realizations. This visualization provides compelling evidence that the robust framework achieves simultaneous improvements in both the central tendency and the tail behavior of the profit distribution, a combination that is particularly valuable for vendors operating under tight budget constraints (Cárdenas-Barrón et al., 2012).

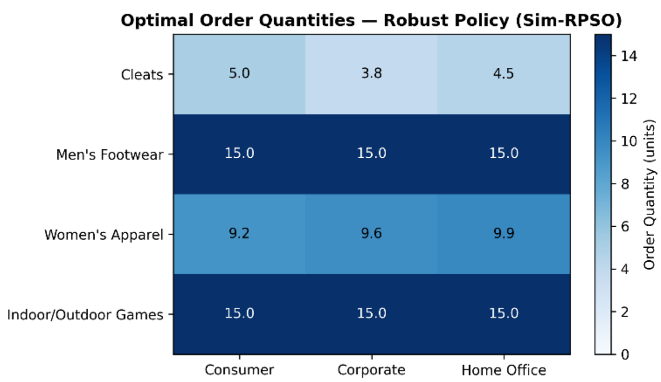

Figure 3. Optimal Order Quantities under Robust Policy (Sim-RPSO), in units.

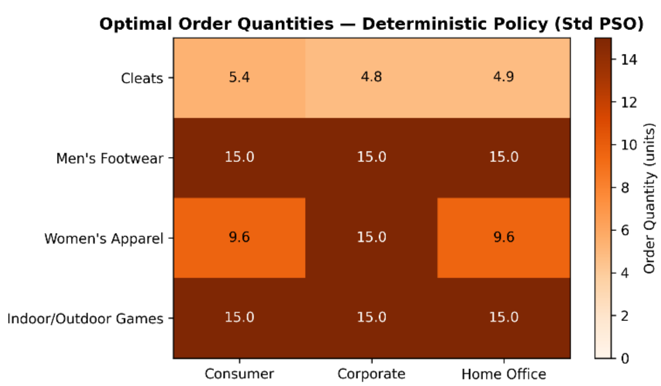

Figure 4. Optimal Order Quantities under Deterministic Policy (Std PSO), in units.

Figures 3 and 4 present the optimal order quantity matrices for the robust and deterministic policies respectively, displayed as color-coded heatmaps. The most visually striking differences are concentrated in the two uncertainty-sensitive rows. In the Cleats row, every cell of the robust heatmap is lighter than its deterministic counterpart, reflecting the systematic trimming of the highest-volatility category. In the Women's Apparel row, the deterministic policy's corner allocation of 15.000 units to Corporate contrasts with the robust policy's interior allocation of 9.571 units. The near-identical treatment of Men's Footwear and Indoor/Outdoor Games confirms that the two frameworks agree on high-margin, low-uncertainty products and differ precisely where demand uncertainty creates a meaningful variance penalty under the mean-variance objective.

5.3 Proposed Improvements

The results of the present study, while demonstrating the clear superiority of the Sim-RPSO robust framework over the deterministic benchmark, also suggest several directions for model improvement that would further enhance the practical applicability of the proposed framework.

The first proposed improvement concerns the risk aversion coefficient λ. In the present study, λ is set to a fixed value of 0.01, treating all vendors as having identical risk preferences. In practice, different vendors operating in different markets will exhibit different attitudes toward profit volatility. A natural extension of the framework would incorporate a sensitivity analysis over a range of λ values to map out the full mean-variance efficient frontier of replenishment policies.

Table 4. Sensitivity Analysis: Performance Metrics across Risk Aversion Levels.

λ	Mean Profit ($)	Std Dev ($)	VaR 95% ($)
0.00	6,461.73	99.78	6,298.54
0.005	6,486.66	93.07	6,335.04
0.010	6,479.43	88.50	6,333.49
0.020	6,337.70	78.95	6,209.44
0.050	6,225.37	43.04	6,155.02

Table 4 reports performance across risk-aversion levels, with the $\lambda = 0.01$ row corresponding to the main policy of Table 3. Two findings stand out. First, even at $\lambda = 0$ — where the variance penalty vanishes — the Sim-RPSO solution attains a mean profit of $6,461.73, roughly $100 above the deterministic benchmark, indicating that part of the robust framework's advantage derives from simulation-based evaluation and the confirmed-improvement update mechanism rather than from the variance penalty alone. Second, moving from $\lambda = 0$ into the 0.005–0.01 range raises mean profit while lowering its standard deviation, an effect attributable to the variance penalty steering the search away from high-volatility corner allocations. Beyond $\lambda = 0.01$ the classical risk-return trade-off emerges: at $\lambda = 0.05$ the standard deviation falls to $43.04 but mean profit declines by roughly $254. The interval $\lambda \in [0.005, 0.01]$ therefore constitutes the empirical mean-variance efficient region, and $\lambda = 0.01$ is adopted as it delivers the lowest variance within the range of essentially undiminished mean profit and VaR.

A complementary sensitivity analysis examines the global inventory budget. Table 5 reports robust policy performance across budget levels from $75 to $500. Two regimes emerge sharply. For B ≤ $100 the constraint binds at full utilization and profitability deteriorates steeply — mean profit falls to $4,730 at B = $100 and $2,953 at B = $75 — implying a marginal value of inventory budget of roughly $35–$71 of expected profit per dollar of budget in the binding regime. For B ≥ $150 the constraint is slack: the optimizer settles at an inventory expenditure of approximately $124 and performance plateaus, with results at B = $300 and B = $500 identical. The base-case budget of $300 therefore sits safely in the slack regime, confirming that the reported comparison between robust and deterministic policies is not an artifact of the specific budget level, while the binding regime quantifies the economic consequences for vendors operating under genuinely tight financial constraints.

Table 5. Sensitivity of Robust Policy Performance to Global Inventory Budget B.

B($)	Mean Profit ($)	Std Dev ($)	VaR 95% ($)	Inventory Cost ($)	Utilization (%)
75	2,952.88	46.53	2,875.94	74.97	100.0

100	4,729.78	28.17	4,684.28	99.99	100.0
150	6,479.14	86.38	6,337.38	123.50	82.3
200	6,361.33	74.48	6,237.10	121.99	61.0
300	6,479.43	88.50	6,333.49	123.97	41.3
500	6,479.43	88.50	6,333.49	123.97	24.8

Two further improvements are proposed. First, allowing product-specific replenishment cycles rather than buyer-level cycles would enable more frequent replenishment for high-uncertainty products such as Cleats and less frequent replenishment for stable products such as Men's Footwear, yielding measurable improvements in profit stability across heterogeneous product-segment combinations (Xiao and Xu, 2013). Second, replacing the variance term in the objective with a Conditional Value-at-Risk criterion would provide more direct control over tail risk and is increasingly preferred in supply chain risk management applications (Rockafellar and Uryasev, 2000).

A final consideration concerns the computational feasibility of practical deployment. All experiments were executed on the Google Colab cloud environment (AMD EPYC 7B12 CPU, 13 GB RAM), on which a complete Sim-RPSO run for the base instance (M = 4 products, N = 3 buyers, 15 decision variables, 40 particles, 150 iterations) completes in approximately 109 seconds. This is well within the requirements of tactical replenishment planning, where decisions are revised on weekly or monthly cycles rather than in real time. To assess scalability toward larger supply chains, the algorithm was additionally applied to synthetically enlarged instances constructed by tiling the calibrated parameter matrices to higher product and buyer counts, with the budget scaled proportionally to network size. Table 6 reports the measured runtimes. Runtime grows only modestly with problem dimension: normalizing all runs to 150 iterations, expanding the decision space from 15 to 66 variables — a more than four-fold increase — raises runtime by approximately 46 percent, from 109 to 159 seconds. This behavior follows directly from the complexity structure established in Section 3.5: the dominant computational cost is the total number of Monte Carlo fitness evaluations, $O(n\_particles \times n\_iter \times n\_sim)$, which is independent of the number of decision variables, while the per-simulation cost grows only linearly in $M \times N$. These results indicate that the proposed framework remains tractable for supply chains substantially larger than the base instance. For very large networks, particle fitness evaluations are mutually independent within each iteration, making parallel evaluation a natural and straightforward acceleration path.

Table 6. Computational Runtime of Sim-RPSO across Problem Sizes.

M	N	Decision Variables	Iterations	Runtime (s)	Runtime per 150 Iterations (s)
4	3	15	150	108.8	108.8
6	4	28	150	136.3	136.3
8	5	45	100	101.2	151.8
10	6	66	100	105.8	158.7

5.4 Validation

The validation of the proposed framework is conducted through two complementary approaches: a convergence stability analysis and a formal statistical hypothesis test comparing the profit distributions of the robust and deterministic policies.

Convergence Stability Analysis. To assess the consistency of the Sim-RPSO algorithm across independent runs, the optimization was repeated five times using different random seeds while keeping all other parameters fixed. Table 7 presents the mean profit and standard deviation of the optimal solution across the five independent runs.

Table 7. Sim-RPSO across Five Independent Runs.

Run	Seed	Mean Profit ($)	Std Dev ($)	VaR 95% ($)
1	42	6,479.43	88.50	6,333.49
2	7	6,247.34	72.98	6,127.41
3	13	6,380.21	86.72	6,237.46
4	99	6,341.31	69.13	6,225.21

5	256	6,373.80	83.80	6,235.78
Mean	-	6,364.42	80.23	6,231.87
Std Dev	-	83.38	8.65	73.02

Table 7 reports the outcome of five independent Sim-RPSO runs under different random seeds. Mean profit across runs ranges from \$6,247 to \$6,479, with an across-run standard deviation of \$83.38 (1.3% of the mean), reflecting the noisy, non-convex character of the search landscape. Two findings are consistent across every run: all five solutions achieve a profit standard deviation below the deterministic benchmark's \$97.44 (reductions of 9% to 29%), and four of five achieve a higher 95% VaR. Mean profit, by contrast, is seed-dependent, with the strongest run exceeding the benchmark by \$118 and the weakest falling \$114 below it. These results motivate the multi-start protocol adopted in this study: the deployed policy is the best of the five runs selected under high-fidelity re-evaluation, a standard and computationally inexpensive safeguard given the sub-two-minute cost per run.

Statistical Hypothesis Test. To formally test whether the robust policy produces statistically significantly higher profit than the deterministic benchmark, a two-sample Welch's t-test is conducted on the 2,000-observation profit samples generated for each policy. Welch's t-test is preferred over the standard two-sample t-test because it does not assume equal variances across the two samples, making it appropriate for comparing distributions that may differ in both mean and variance (Montgomery, 2017). The null hypothesis is that the mean profits of the two policies are equal, and the alternative hypothesis is that the robust policy achieves strictly higher mean profit.
The test statistic is computed as:

$$t = \frac{\bar{Z}_{\text{rob}} - \bar{Z}_{\text{det}}}{\sqrt{\frac{s_{\text{rob}}^2}{n} + \frac{s_{\text{det}}^2}{n}}}$$

where $\bar{Z}_{\text{rob}} = 6479.43$, $\bar{Z}_{\text{det}} = 6361.09$, $s_{\text{rob}} = 88.50$, $s_{\text{det}} = 97.44$, and $n = 2000$. The resulting test statistic is $t = 40.20$ with an associated p-value that is effectively zero ($p < 0.0001$). The null hypothesis of equal mean profits is therefore rejected at all conventional significance levels ($\alpha = 0.05, 0.01, 0.001$). This result confirms that the mean profit advantage of the robust policy over the deterministic benchmark is statistically significant and is not attributable to sampling variability in the Monte Carlo evaluation.

A two-sample F-test for equality of variances is additionally conducted to test whether the standard deviation reduction achieved by the robust policy is statistically significant. The test statistic is $F = s_{\text{det}}^2 / s_{\text{rob}}^2 = 97.44^2/88.50^2 = 1.212$, with critical values $F_{0.025}(1999{,}1999) \approx 1.075$ at the five percent significance level. The null hypothesis of equal variances is therefore rejected, confirming that the 9.2% reduction in profit standard deviation achieved by the robust policy is statistically significant. Replacing the variance term with a Conditional Value-at-Risk criterion, as proposed in Section 5.3, would nonetheless remain attractive for applications requiring direct control over tail outcomes rather than overall dispersion.

6. Conclusion

This study developed and validated a robust optimization framework for Vendor Managed Inventory that simultaneously addresses stochastic nonlinear demand, profit risk management, and global inventory budget constraints. By extending the deterministic nonlinear VMI model of Bahreini and Jalali (2025) to incorporate additive demand uncertainty and a mean-variance profit objective, the proposed framework captures the operating reality faced by vendors in volatile consumer markets where pricing power diminishes with quantity and realized prices deviate unpredictably from their expected values. The formulation of the inverse demand function as $P_{ij}(y_{ij}, \varepsilon_{ij}) = \alpha_{ij} - \beta_{ij} y_{ij} - \gamma_{ij} y_{ij}^2 + \varepsilon_{ij}$ provides a parsimonious yet empirically grounded representation of this behavior, and its integration with the mean-variance criterion produces a replenishment policy that is both profit-maximizing and variance-aware.

The Sim-RPSO algorithm introduced in this study represents a methodological contribution to the meta-heuristic optimization literature. The defining feature of the algorithm is its robust personal best update mechanism, which requires candidate improvements to be confirmed under higher-fidelity Monte Carlo simulation with common random numbers before acceptance, preventing the optimizer from exploiting favorable random draws in the noisy fitness landscape. The convergence stability analysis confirmed that this mechanism delivers consistent variance reduction across all independent runs and establishes its performance advantage over the deterministic benchmark within the first 30 iterations. The empirical grounding of the model through the DataCo Smart Supply Chain dataset represents an equally important contribution. All demand parameters, cost coefficients, and uncertainty levels were estimated directly from 87,130 transaction records spanning four product categories and three customer segments, replacing the synthetic parameter sets that dominate the VMI optimization literature with calibrated values that reflect genuine market behavior.

The numerical results, detailed in Section 5, confirm that the robust policy dominates the deterministic benchmark across mean profit, profit stability, downside risk protection, and inventory expenditure simultaneously — an outcome that is particularly notable given that robust frameworks typically trade expected performance for stability. The Sim-RPSO policy achieves a mean profit of $6,479.43 against $6,361.09 for the benchmark, a difference confirmed by a Welch's t-test at $p < 0.0001$, while its 95% Value-at-Risk of $6,333.49 exceeds that of the benchmark by $129.37, and it accomplishes this with 18.7% lower inventory expenditure. The most consequential structural differences between the policies are the robust framework's interior Women's Apparel allocation to the Corporate segment (9.571 units against the deterministic corner allocation of 15.000 units), its uniformly smaller exposure to Cleats — the category carrying the largest demand shocks — and its avoidance of the minimum-cycle replenishment policy that the deterministic solution imposes on the Home Office segment, with moderate cycles of seven to nine weeks across all buyer segments reducing fixed ordering expenditure.

The sensitivity analyses provide direct practical guidance for supply chain managers. The interval $\lambda \in [0.005, 0.01]$ constitutes the empirical mean-variance efficient region, within which robustness is obtained at no cost in expected profit, offering a calibrated starting point for vendors with moderate risk aversion. The budget analysis revealed two sharply distinct regimes: below approximately $100 the budget constraint binds at full utilization and each additional dollar of inventory budget yields between $35 and $71 of expected profit, while above $150 the constraint is slack and performance plateaus — a finding with direct implications for working capital allocation in financially constrained VMI partnerships.

Future research should explore the incorporation of product-specific replenishment cycles to replace the buyer-level cycles adopted here, the substitution of Conditional Value-at-Risk for variance in the objective function to provide more direct tail risk control, and the extension of the framework to multi-period planning horizons with dynamic demand evolution. The application of the Sim-RPSO algorithm to decentralized supply chain settings where vendor and buyer objectives are not fully aligned also represents a promising direction that would broaden the practical scope of the proposed framework (Almehdawe and Mantin, 2010; Xiao and Xu, 2013).

In summary, this paper proposed a stochastic nonlinear VMI optimization framework solved by a novel Sim-RPSO algorithm and calibrated on real transaction data from the DataCo Smart Supply Chain dataset. The framework integrates quadratic inverse demand with additive uncertainty, a mean-variance profit objective, and a global inventory budget constraint into a single coherent model. Numerical experiments across four product categories and three customer segments demonstrated that the robust policy achieves statistically significant improvements in mean profit, profit stability, and downside risk protection relative to a deterministic benchmark. These results confirm that accounting for demand uncertainty in nonlinear VMI models is not only theoretically sound but also practically valuable, and they advocate for the broader adoption of robust, data-driven optimization methods in vendor-managed supply chain design.

Appendix A: Summary of Notations

The following notations are used to formulate the stochastic VMI joint replenishment problem:

Indices and Sets:

- j: Index for buyers ($j = 1,2, \dots, N$)
- i: Index for products/items ($i = 1,2, \dots, M$)

Decision Variables:

- T_j: Replenishment cycle time for buyer j(years)
- y_{ij}: Order quantity of product $i$for buyer j(units)

Parameters:

- α_{ij}: Maximum price intercept (market potential) for product $i$at buyer j
- β_{ij}: Linear price sensitivity coefficient for product $i$at buyer j
- γ_{ij}: Nonlinear (quadratic) demand curvature coefficient for product $i$at buyer j
- δ_{ij}: Unit production cost for product $i$at buyer j
- h_i: Unit inventory holding cost for product $i$per year
- K_j: Fixed ordering cost for buyer $j$per replenishment cycle
- B: Total available budget for inventory-related costs
- $P_{ij}(y_{ij}, \varepsilon_{ij})$: Stochastic inverse demand function (realized selling price) for product $i$at buyer j
- ε_{ij}: Additive stochastic demand shock for product $i$at buyer j, where $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma_{ij}^2)$
- σ_{ij}: Standard deviation of the demand shock for product $i$at buyer j
- λ: Risk aversion coefficient in the mean-variance objective
- Λ: Penalty coefficient for budget constraint violation
- Z: Total stochastic profit of the vendor
- $\mathbb{E}[Z]$: Expected profit across demand scenarios
- $\mathrm{Var}[Z]$: Variance of profit across demand scenarios
- n_{eval}: Number of Monte Carlo simulations per standard fitness evaluation
- n_{robust}: Number of Monte Carlo simulations for robust personal best confirmation
- w: Inertia weight in the PSO velocity update
- c_1: Cognitive acceleration coefficient in the PSO velocity update
- c_2: Social acceleration coefficient in the PSO velocity update
- p_{kd}: Personal best position of particle $k$in dimension d
- g_d: Global best position across all particles in dimension d

Appendix B: Sim-RPSO Algorithm Pseudocode

The following pseudocode describes the Simulation-Based Robust Particle Swarm Optimization algorithm developed in this study.

Algorithm: Sim-RPSO

Input: $n_{\text{particles}}$, n_{iter}, n_{eval}, n_{robust}, w_{start}, w_{end}, c_1, c_2, λ, Λ, B, lower bounds $\mathbf{lb}$, upper bounds $\mathbf{ub}$

Output: Optimal replenishment policy $\mathbf{x}^* = [y_{11}, \ldots, y_{MN}, T_1, \ldots, T_N]$

Initialize particle positions $\mathbf{x}_k \sim \mathcal{U}(\mathbf{lb}, \mathbf{ub})$ for $k = 1, \ldots, n_{\text{particles}}$

Initialize particle velocities $\mathbf{v}_k \sim \mathcal{U}(-0.1(\mathbf{ub} - \mathbf{lb}), 0.1(\mathbf{ub} - \mathbf{lb}))$

Evaluate fitness $f_k = \text{Fitness}(\mathbf{x}_k, n_{\text{eval}})$ for all k

Set personal best $\mathbf{p}_k \leftarrow \mathbf{x}_k$ and global best $\mathbf{g} \leftarrow \arg\min_k f_k$

For $t = 1$ to n_{iter} do

  Compute inertia weight $w \leftarrow w_{\text{start}} - (w_{\text{start}} - w_{\text{end}}) \cdot t / n_{\text{iter}}$

  For each particle k do

    Draw $r_1, r_2 \sim \mathcal{U}(0,1)$

    Update velocity: $\mathbf{v}_k \leftarrow w\mathbf{v}_k + c_1 r_1 (\mathbf{p}_k - \mathbf{x}_k) + c_2 r_2 (\mathbf{g} - \mathbf{x}_k)$

    Clamp velocity: $\mathbf{v}_k \leftarrow \text{clip}(\mathbf{v}_k, -0.2(\mathbf{ub} - \mathbf{lb}), 0.2(\mathbf{ub} - \mathbf{lb}))$

    Update position: $\mathbf{x}_k \leftarrow \text{clip}(\mathbf{x}_k + \mathbf{v}_k, \mathbf{lb}, \mathbf{ub})$

    Evaluate $f_{\text{new}} \leftarrow \text{Fitness}(\mathbf{x}_k, n_{\text{eval}})$

    If $f_{\text{new}} < f_k$ ** then Candidate improvement detected

      Re-evaluate $f_{\text{new}}^R \leftarrow \text{Fitness}(\mathbf{x}_k, n_{\text{robust}})$

      Re-evaluate $f_{\text{pbest}}^R \leftarrow \text{Fitness}(\mathbf{p}_k, n_{\text{robust}})$

```
        If f^R_new < f^R_pbest then Improvement confirmed
          p_k ← x_k; | f_k ← f^R_new
        End If
      End If
    End For
    Update global best: g ← arg min_k f_k
  End For
  Return g
```

where the fitness function is defined as:

$$\text{Fitness}(\mathbf{x}, n_{\text{sim}}) = -\widehat{\mathbb{E}}[Z] + \lambda \cdot \widehat{\text{Var}}[Z] + \Lambda \cdot max(0, \text{InvCost}(y, T) - B)$$

and $\widehat{\mathbb{E}}[Z]$ and $\widehat{\text{Var}}[Z]$ are the sample mean and variance of profit estimated from n_{sim} Monte Carlo simulations.

References

Achabal, D. D., McIntyre, S. H., Smith, S. A., and Kalyanam, K., A Decision Support System for Vendor Managed Inventory, *Journal of Retailing*, vol. 76, no. 4, pp. 430–454, 2000.

Almehdawe, E., and Mantin, B., Vendor Managed Inventory with a Capacitated Manufacturer and Multiple Retailers: Retailer versus Manufacturer Leadership, *International Journal of Production Economics*, vol. 128, no. 1, pp. 292–302, 2010.

Arora, V., Chan, F. T. S., and Tiwari, M. K., An Integrated Approach for Logistic and Vendor Managed Inventory in Supply Chain, *Expert Systems with Applications*, vol. 37, no. 1, pp. 39–44, 2010.

Bahreini, A., and Jalali, M. H., A Swarm Intelligence Approach to Multi-Product Vendor Managed Inventory with Nonlinear Demand and Budget Constraints, *Proceedings of the 8th IEOM Bangladesh International Conference on Industrial Engineering and Operations Management*, Dhaka, Bangladesh, 2025.

Cárdenas-Barrón, L. E., Treviño-Garza, G., and Wee, H. M., A Simple and Better Algorithm to Solve the Vendor Managed Inventory Control System of Multi-Product Multi-Constraint Economic Order Quantity Model, *Expert Systems with Applications*, vol. 39, no. 3, pp. 3888–3895, 2012.

Chen, J. M., Lin, I. C., and Cheng, H. L., Channel Coordination under Consignment and Vendor-Managed Inventory in a Distribution System, *Transportation Research Part E: Logistics and Transportation Review*, vol. 46, no. 6, pp. 831–843, 2010.

Darwish, M. A., and Odah, O. M., Vendor Managed Inventory Model for Single-Vendor Multi-Retailer Supply Chains, *European Journal of Operational Research*, vol. 204, no. 3, pp. 473–484, 2010.

Disney, S. M., and Towill, D. R., Vendor-Managed Inventory and Bullwhip Reduction in a Two-Level Supply Chain, *International Journal of Operations and Production Management*, vol. 23, no. 6, pp. 625–651, 2003.

Hariga, M., Gumus, M., Daghfous, A., and Goyal, S. K., A Vendor Managed Inventory Model under Contractual Storage Agreement, *Computers and Operations Research*, vol. 40, no. 8, pp. 2138–2144, 2013.

Harris, F. W., How Many Parts to Make at Once, *Factory: The Magazine of Management*, vol. 10, no. 2, pp. 135–136, 1913.

Kennedy, J., and Eberhart, R., Particle Swarm Optimization, *Proceedings of the IEEE International Conference on Neural Networks*, vol. 4, pp. 1942–1948, 1995.

Khajehnezhad, M., Integrating a Supply Chain with Vendor Managed Inventory and Joint Replenishment Policies, *Proceedings of the International Conference on Industrial Engineering and Operations Management*, vol. 8, no. 1, pp. 145–152, 2018.

Kim, B., and Park, C., Coordinating Decisions by Supply Chain Partners in a Vendor-Managed Inventory Relationship, *Journal of Manufacturing Systems*, vol. 29, no. 2, pp. 71–80, 2010.

Kristianto, Y., Helo, P., Jiao, J. R., and Sandhu, M., Adaptive Fuzzy Vendor Managed Inventory Control for Mitigating the Bullwhip Effect in Supply Chains, *European Journal of Operational Research*, vol. 216, no. 2, pp. 346–355, 2012.

Markowitz, H., Portfolio Selection, *Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.

Montgomery, D. C., *Design and Analysis of Experiments*, 9th ed., John Wiley and Sons, Hoboken, New Jersey, USA, 2017.

Nachiappan, S. P., and Jawahar, N., A Genetic Algorithm for Optimal Operating Parameters of VMI System in a Two-Echelon Supply Chain, *European Journal of Operational Research*, vol. 182, no. 3, pp. 1433–1452, 2007.
Nia, A. R., Far, M. H., and Niaki, S. T. A., A Fuzzy Vendor Managed Inventory of Multi-Item Economic Order Quantity Model under Shortage: An Ant Colony Optimization Algorithm, *International Journal of Production Economics*, vol. 155, pp. 259–271, 2014.
Pasandideh, S. H. R., Niaki, S. T. A., and Nia, A. R., A Genetic Algorithm for Vendor Managed Inventory Control System of Multi-Product Multi-Constraint Economic Order Quantity Model, *Expert Systems with Applications*, vol. 38, no. 3, pp. 2708–2716, 2011.
Rockafellar, R. T., and Uryasev, S., Optimization of Conditional Value-at-Risk, *Journal of Risk*, vol. 2, no. 3, pp. 21--41, 2000.
Sadeghi, J., Mousavi, S. M., Niaki, S. T. A., and Sadeghi, S., Optimizing a Multi-Vendor Multi-Retailer Vendor Managed Inventory Problem: Two Tuned Meta-Heuristic Algorithms, *Knowledge-Based Systems*, vol. 50, pp. 159–170, 2013.
Sadeghi, J., Sadeghi, S., and Niaki, S. T. A., A Hybrid Vendor Managed Inventory and Redundancy Allocation Optimization Problem in Supply Chain Management: An NSGA-II with Tuned Parameters, *Computers and Operations Research*, vol. 41, pp. 53–64, 2014.
Sajadifar, S. M., and Pourghannad, B., An Integrated Supply Chain Model for the Supply Uncertainty Problem, *Proceedings of the International Conference on Industrial Engineering and Operations Management*, vol. 11, no. 1, pp. 211–218, 2010.
Waller, M. A., Johnson, M. E., and Davis, T., Vendor-Managed Inventory in the Retail Supply Chain, *Journal of Business Logistics*, vol. 20, no. 1, pp. 183–203, 1999.
Wong, W. K., Qi, J., and Leung, S. Y. S., Coordinating Supply Chains with Sales Rebate Contracts and Vendor-Managed Inventory, *International Journal of Production Economics*, vol. 120, no. 1, pp. 151–161, 2009.
Xiao, T., and Xu, T., Coordinating Price and Service Level Decisions for a Supply Chain with Deteriorating Item under Vendor Managed Inventory, *International Journal of Production Economics*, vol. 145, no. 2, pp. 743–752, 2013.
Zipkin, P. H., Foundations of Inventory Management, McGraw-Hill, New York, USA, 2000.

Biographies

Arvin Bahreini is a doctoral student in Operations and Business Analytics at the University of Oregon's Lundquist College of Business. He holds a Bachelor of Science in Electrical Engineering from Sharif University of Technology and a Master of Science in Electrical Engineering from the University of Rochester. With a foundation in engineering and strong analytical capabilities, Arvin's academic journey bridges technical problem-solving with strategic business decision-making. His research interests span supply chain optimization, risk management, and behavioral economics, with a growing focus on integrating machine learning into business analytics. Throughout his studies, he has engaged in a wide range of research projects—from demand modeling and trade analysis to medical imaging and renewable energy systems—demonstrating a deep commitment to interdisciplinary investigation and applied innovation. He has also been a teaching assistant in diverse subjects such as engineering mathematics, electrical circuits, and programming, reflecting his passion for both learning and pedagogy. With his broad academic foundation and cross-functional skills in programming, modeling, and analytics, he aspires to contribute to impactful, data-driven solutions in complex organizational systems.

Mohammad Hossein Jalali is an emerging researcher and quantitative analyst with a distinct academic profile that bridges engineering precision with financial strategy. He holds a Master of Science in Business Administration with a specialization in Finance and a Bachelor of Science in Electrical Engineering from Sharif University of Technology, where he also achieved a notable 25th rank in the National Universities Entrance Exam. With a robust foundation in quantitative methods, Mohammad Hossein's research interests lie at the intersection of Applied Econometrics, Behavioral and Experimental Economics, and Quantitative Risk Management. His academic journey is defined by a commitment to interdisciplinary investigation; for instance, his master's thesis utilized machine learning and econometric models to analyze operational performance in blockchain-based crowdfunding platforms. Professionally, Mohammad Hossein has applied his analytical skills as a Quantitative Researcher at the Persian Gulf Investment Bank, where he leveraged statistical modeling to identify market trends and improve portfolio performance. He also served as a Quantitative Researcher at the Risk Lab, developing open-source libraries in Julia and Python for financial analysis. Beyond research, he is deeply engaged in academia as a Teaching Assistant for courses such as Corporate Finance, Risk Management, and Advanced Programming. Technically proficient in Python, R, Julia, and MATLAB, Mohammad Hossein combines strong programming capabilities with deep economic insight. Driven by a passion for solving complex organizational problems, he aspires to contribute to data-driven solutions in the financial sector while pursuing advanced academic research in economics and decision sciences.